Security Against Compelled Disclosure

Ian Brown
Hidden Footprints Ltd.
6 St. Andrew's St.
London EC4A 3LX
United Kingdom
ianb@hfprints.com

Ben Laurie
A.L. Digital Ltd.
Voysey House
Barley Mow Passage
London W4 4GB
United Kingdom
ben@algroup.co.uk

Abstract

Various existing and pending legislation can be used to force individuals and organisations to disclose confidential information. Courts may order a wide variety of data to be turned over by either party in civil and criminal cases. Government agencies are explicitly tasked with protecting "national economic security." And organised crime will target information just like any other valuable asset. In a less than perfectly ethical world, companies require means to protect their information assets against economic espionage, misuse of discovery processes and criminal coercion. We describe actual and potential examples of compelled disclosure abuses in the US and UK, and legal enhancements to conventional security services for protecting communications and stored data against their recurrence.

1. Introduction

Traditional security threat models assume a powerful adversary, with access to all communications links and insecure data and systems. Hugely- funded intelligence organisations such as the US National Security Agency come close to filling this role. But with powers to compel disclosure of virtually any information, governments and courts present an even greater challenge to security system designers.

Whilst it would be unethical to create defences against these powers if they were always correctly used, the combination of increasing government interest in economic intelligence and the multinational nature of more and more organisations inevitably encourages government agencies to consider all available mechanisms to provide economic intelligence to their customers.

Valuable information will also increasingly become the target of criminal activity. Data is already the most important resource of many organisations.

This paper reviews the threats to information assets presented by the misuse of the judicial discovery process, Customs powers, newly emerging key disclosure warrants, and misdirected signals intelligence product. Old-fashioned intimidation and blackmail are also considered.

Companies are vulnerable not just to economic intelligence gathering, but breach of due care obligations such as under non-disclosure agreements, revelation of trade secrets, or even the exposure of the identity of whistleblowers. Compromise of cryptographic keys can seriously damage an organisation's security for significant periods of time.

We then describe legal enhancements to conventional security services that increase the resistance of communications and data security measures to these threats, and procedural safeguards in their use. Governments already fully realise the need for some of these measures, which are widely deployed in military systems such as the US Defence Messaging System.

Finally we discuss the information ethics required of governments to reduce the need for these efforts.

2. Current and pending disclosure legislation

2.1. Judicial discovery processes

Courts the world over have extensive powers to order the production of information by parties to a case. Witnesses can be compelled to give evidence and produce documents in intelligible form. Non-compliance can lead to criminal contempt of court charges, or loss of the case. US litigation in particular relies heavily on a pre-trial discovery process where the parties may judge the strength of their case by viewing their opponents' evidence.

The anti-trust case against Microsoft has shown how devastating these powers can be. The initial Justice Department

complaint relied heavily on internal Microsoft e-mail obtained during their investigation. Executive after executive, all the way up to Bill Gates, were extensively quoted describing how Windows should be leveraged to beat Netscape in the browser wars. Microsoft's Christian Wildfeuer, for example, allegedly wrote that "It seems clear that it will be very hard to increase browser market share on the merits of [Internet Explorer] 4 alone. It will be more important to leverage the OS asset to make people use IE instead of Navigator" [23]. Faced with such evidence, it seems unsurprising that Judge Thomas Penfield Jackson eventually found that Microsoft had abused their dominant position in the operating system market. These types of comments should simply never be put in permanent form.

The Fourth and Fifth Amendments to the US Constitution provide limited protection against search and seizure and compelled self-incrimination for individuals. But a series of Supreme Court judgments have made this protection limited indeed in the case of information existing in physical form [38]. Company documents are explicitly exempt: "The [Fifth] amendment is limited to a person who shall be compelled in any criminal case to be a witness against himself; and if he cannot set up the privilege of a third person, he certainly cannot set up the privilege of a corporation" [5]. And only when a cryptographic key is memorised rather than written down is an individual protected if its disclosure would provide 'testimonial' evidence of criminal activity [38].

Groups have used these processes to discover, among other things, the identity of pseudonymous on-line critics. The Church of Scientology subpoened the e-mail address behind a pseudonym at `penet.fi`, one of the original anonymising remailers. Unfortunately for them, it pointed only to another pseudonym at a more advanced remailer at `c2.net`. ITEX Corporation sued Yahoo to discover the identity of 100 "John Doe" critics who had made negative comments on a Yahoo Finance message board [26]. In neither case were the individuals behind the pseudonyms given the opportunity to first present the case against their exposure.

2.2. Import and export searches

Customs authorities have very wide-ranging powers to search materials being imported to and exported from a country. They have recently used these powers to investigate digital information such as that on laptop hard disks. UK Customs and Excise have declared that officials will routinely scan laptops for illegal material such as pornography [12] without any requirement for probable cause as exists in the US. Journalist Ken Cukier reported that Customs randomly attempted to scan his laptop on arriving in London on Eurostar (but were thwarted due to his use of an Apple machine) [30]. Whether scans will also be used to further the "economic well-being" of the UK must be judged by those carrying information that would help that aim. W. H. Murray, an information security consultant with Deloitte and Touche, concluded:

"While I may not have anything on my laptop that C & E have any legitimate objection to, I have much on my laptop that is none of HM's government's legitimate business. Some of it is personal. Much of it is data of or about my clients which I have a professional, ethical, and contractual obligation to keep confidential. Some of it might never have been shared if such procedures had been routine, to my detriment, to that of my clients, and to that of the commonwealth... I will have to advise my clients to eschew discretionary travel through or to a country that has such procedures. Since I never travel without my computer, would not be much value to my clients without it, and carry on it information which I owe a duty to my clients to protect from copying, I will not go back to England unless and until HM's government renounces such extreme measures. As it is one of my favorite destinations, I do not say that lightly." [29]

2.3. Decryption and key warrants

The Organisation for Economic Cooperation and Development guidelines on encryption state that "National cryptography policies may allow lawful access to plaintext, or cryptographic keys, of encrypted data." [31] A number of governments have considered legislation that would require plaintext or keys to be provided by an individual or organisation under warrant. So far the UK, India, Singapore and Malaysia have implemented such legislation. The US, Belgium and Netherlands have acts pending that require third parties in possession of keys to provide them to authorities [4]. The draft Council of Europe Convention on Cyber-Crime [11] requires signatories to provide legislative powers to force individuals to reveal any reasonable information required to search or copy seized secure data (cl.2(5)).

The Regulation of Investigatory Powers Act [33] provides these powers in the UK. Its decryption or key request warrants may also contain a gagging clause that carries five years' imprisonment for an individual notifying anyone other than their lawyer that the warrant has been served. The UK Home Secretary may authorise warrants "for the purpose of safeguarding the economic well-being of the United Kingdom" (s.5(3)c). And warrants may be served "for the purpose of securing the effective exercise or proper performance by any public authority of any statutory power or statutory duty" (s.49(2)bii) — a *huge* number of bodies in the UK. Serious doubts have been raised about the ability of public authorities to provide security for disclosed keys commensurate with their value [21].

Ireland's Electronic Commerce Act 2000 [15] requires

that seized encrypted information be put into intelligible form (s.26), but explicitly states that "Nothing in this Act shall be construed as requiring the disclosure or enabling the seizure of unique data, such as codes, passwords, algorithms, private cryptographic keys, or other data, that may be necessary to render information or an electronic communication intelligible" (s.27). While this still leaves documents vulnerable, it prevents the long-term damage to company security systems caused by key compromise.

2.4. Signals Intelligence

Perhaps the most pervasive threat to corporate information is the activities of the world's Signals Intelligence (SIGINT) agencies. These government organisations use a vast array of technologies to capture communications from commercial satellites, long distance communications, undersea cables, and at many points on the Internet. More than 120 satellites are in operation to support their activities. Members of the five-nation UKUSA alliance (the US, UK, Canada, Australia and New Zealand) share SIGINT facilities, tasks and product [7].

While the National Security Agency, Britain's Government Communications Headquarters, and their many foreign equivalents have become slightly better known during the past twenty years, few outside the security community realise the vast scale of their activities. The 1998 NSA budget is estimated at $3.6bn, with significant further costs incurred in the $6.3bn National Reconnaisance Office expenditures [32]. Globally, it is estimated that $15–20bn is spent every year on communications intelligence [7].

The US Foreign Intelligence Advisory Board recommended in 1970 that "henceforth economic intelligence be considered a function of the national security, enjoying a priority equivalent to diplomatic, military, technological intelligence" [6]. The National Security Agency is authorised by Executive Order 12333 to collect "information for national foreign intelligence purposes in accordance with guidance from the Director of Central Intelligence" (s.1.12(b)3) [17]. The secret Office of Intelligence Liason, renamed in 1993 to the Office of Executive Support, routes intelligence information to the Department of Commerce, from where "tips based on spying ... regularly flow from the Commerce Department to U.S. companies to help them win contracts overseas" [39]. A report from the European Parliament states that deals worth billions and billions of dollars have been won by US companies after receiving signals intelligence from their government [7].

While claiming US economic espionage is limited to preventing bribery winning contracts, ex-Director of Central Intelligence James Woolsey admitted that "We steal [economic] secrets with espionage, with communications, with reconnaissance satellites" [44]. And letters from the CIA to Congress disclose that the intelligence agencies are also interested in "lobbying," "linking financial aid to contract awards" and "the use of insider information and disinformation against U.S. firms." The CIA's National Counter Intelligence Center reported to Congress that "because they are so easily accessed and intercepted, corporate telecommunications — particularly international telecommunications — provide a highly vulnerable and lucrative source for anyone interested in obtaining trade secrets or competitive information." [43]

GCHQ provides similar economic intelligence in the UK. It is authorised by the Intelligence Services Act 1994 [24] to intercept foreign communications "in the interests of the economic well-being of the United Kingdom" (s.3(2)b). Targets may be specified by the government's Overseas Economic Intelligence Committee, the Economic Section of the Joint Intelligence Committee, the Treasury, or the Bank of England [42]. MI5, Britain's internal Security Service, is spending £25m on a Government Technical Assistance Centre at its London headquarters to monitor traffic and attempt to decrypt captured ciphertext [35].

2.5. Criminal coercion

The oldest threat against security systems is coercion of authorised users. The strongest cipher is no protection if a key can be obtained from its owner. Bank vaults require two separate keyholders for opening as much to protect bank managers' families against kidnapping as to protect against corrupt staff.

System designers must remember the vulnerabilities introduced by users who may have guns held to their heads, be entrapped by sexually attractive enemy agents, or be subject to a myriad of attacks dreamt up by imaginative criminals and spy thriller authors!

3. Safeguarding communications

All communications over insecure links such as the Internet should be encrypted to protect their contents. But subtle differences in the way secure communication protocols use keys have important ramifications for the security of the resulting ciphertext against disclosure threats.

Online systems such as IPSEC [2] can negotiate new keys for every communication using a protocol such as Diffie-Hellman [13]. If a key is compromised, only the specific session it protected will be revealed to an attacker. This desirable property is called forward secrecy. The security of previous or future encrypted sessions is not affected. Keys are securely deleted after use [14]. Without these keys, there is no way captured ciphertext can be decrypted. The private signature keys used to authenticate online sessions are normally exempt from key disclosure warrants, but can anyway

only be used to intercept a future session using an active attack rather than decipher previously captured ciphertext.

It is more difficult to make store and forward systems like e-mail forward secret, as they rarely make direct connections between a message sender and its recipient. In a typical e-mail encryption system, users create a long-term key pair using a cryptosystem such as RSA [34] and publish the public key on their Web page, in a directory, or via other methods. While the use of long-term keys reduces the administrative burden of key distribution, the practice introduces vulnerabilities. If a public key is used for several years, as is common with systems such as PGP [45], compromise of the private key will allow an attacker to decrypt any message captured during that time. It is therefore important that confidentiality keys are changed regularly. This section describes key management procedures that minimise the lifetime of keys.

It almost goes without saying that *all* communications data should be encrypted. Using security services only for sensitive traffic will immediately arouse suspicion in the data, its sender and recipient, and provide unnecessary information to an attacker performing traffic analysis. We also describe further steps to protect traffic data, such as using an anonymous IP network.

3.1. Short-lifetime encryption keys

Many offline encryption protocols use a public-key cryptosystem to bootstrap a symmetric encryption cipher. The symmetric protocol uses a random key which is then encrypted using the public key of the recipient. The actual content is encrypted using this symmetric key (often referred to as a session key). Three pieces of information are therefore vulnerable to disclosure:

- The private key
- The session key
- The plaintext of the communication

Disclosure of the private key will make all information protected under it vulnerable. Disclosure of a session key or plaintext are equivalent and compromise only that specific session. Britain's RIP Act [33] explicitly allows the seizure of private keys.

Using a series of private keys, each with a short lifetime, reduces the information revealed by the compromise of any one private key because each key protects less data. A long-term signing key is used to certify a succession of short lifetime keys. Each short lifetime key is destroyed soon after its validity period ends (after allowing time for messages in transit to arrive.)

If a message sender has more than one valid encryption key available for a user, they should use the key nearest its expiration date. This limits the time during which the corresponding private key will be available to an attacker.

Key distribution can be eased by submitting new keys to key directories, where they will be available for other users to retrieve. Submission and retrieval can be performed automatically by software. Expired public encryption keys can be deleted by users and keyservers to save space.

PGP allows a user's public key certificate to contain one top-level signature key, which can be used to verify the signatures on further public encryption subkeys, all with different lifetimes. This allows a user to implement as fine-grained short lifetime encryption keys as their security policy requires.

This behaviour is no more difficult to implement in theory with X.509 [9] certificates, since they already have the structure of a chain/hierarchy of signing certificates leading to an encryption certificate, which does not need to be known in advance. However, the current business model of CAs and existing software do not support this mode of operation particularly well. There is also a problem with online retrieval of certificates, a facility not yet widely supported.

3.2. One-time keys

Taking short-term keys to their logical conclusion, a different key could be used to protect every message. Schneier and Hall [36] suggested a user could make several public keys available in a directory. After a key was retrieved by another user, it would be deleted. This requires message senders to have online access to a directory. Not all e-mail users have this facility. It also necessitates trust in the directory operator to act correctly, and is open to a trivial denial of service attack.

An off-line scheme is more compatible with the store and forward nature of e-mail. Every time a user sends a message, they can include a new public key for the recipient to encrypt any reply with.

Users would still possess a relatively long-lived encryption key. If Alice were writing to Bob for the first time, she would encrypt her message with his long term key. She would also include a newly created public key. Bob would use this new key the next time he wrote to Alice, and Alice would decrypt with the associated private key. When Alice sent Bob a new public key and received a reply encrypted with it, she would securely delete the previous private key.

This scheme would use each key pair exactly once if correspondants communicated in sequence. If Bob wanted to send Alice another message before receiving her reply, he would need to use the same public key. Compromise of the associated private key would therefore allow an attacker to decrypt more than one message. If Alice anticipates this situation, she may therefore wish to include a different public key for every reply she expects to that message.

If Alice writes to a large number of people only once or twice, she will generate a large number of private keys that will never be needed. One-time keys should therefore have a short lifetime. As a user's collection of private keys grows, she may wish to reduce the lifetime of new keys.

Note that it would be a sensible precaution to say nothing more in the first message than "please send me a new key". In fact, with increasingly large numbers of people either permanently online, or accessing email frequently, it may make sense to semi-automate this operation.

Users may become frustrated if they are forced to wait every time new keys are generated. Software could prevent this using background key generation, or trade-offs that speed up key generation with minimal reduction in security. With Elgamal [16], for example, the expensive key component to generate is the public prime modulus. A group of keys can share a common public modulus with no negative security implications other than that the key then presents a fatter target for pre-computation attacks.

3.3. Traffic analysis

Even when the contents of a communication are fully secure, its mere existence, timing and recipients can be very valuable information. Such traffic data can reveal previously confidential links between companies, illuminate information from open sources such as reports of a potential takeover bid, and much else besides.

James Bamford has described the importance of traffic analysis to NSA. "Should the cryppies run short on sorcery, the traffic analysts may still be able to salvage a sizable chunk of intelligence. Working only with the 'externals' of the message — where it came from, its apparent destination, the priority, grade of cipher system used, as well as the frequency and volume of other messages — the traffic analysts can often supply the missing piece of a much larger puzzle. A sharp increase in traffic to and from Tyuratam, for example, may indicate an imminent space launch; a sudden switch into a high-grade cipher system or unusual jump in priority by units stationed along the border with Afghanistan may mean an outbreak of hostilities." [3]

Further information is often unwittingly revealed by secure communications software. Secure e-mail programs like PGP can encrypt a message body, but leave header information — crucially, the subject line — as cleartext. Remedies have been proposed that encapsulate the entire message inside another encrypted mail [18], but are as yet unimplemented.

The open architecture of the Internet provides many points at which traffic data may be gleaned from even a fully encrypted connection. But services such as Onion Routing [22] based on Chaumian mixes [10] are now developing to protect this information. These anonymising IP servers receive a multiply-encrypted packet, remove one layer of encryption, and forward the result either to another anonymising server or its final destination. Clients construct a route through several different servers through which their packets travel. Every server on that route must be compromised to reveal a packet's destination. Application-layer data can be sanitised of identifying information so that even the recipient cannot determine the source.

Combining Onion Routing with continuously generated dummy traffic (which is indistinguishable in transit from real traffic, since it is encrypted) and traffic shaping should defeat even the most sophisticated traffic analysis. This comes at a high bandwidth cost, however. Dynamic traffic masking varies the amount of padding traffic to increase efficiency, but without care statistical techniques can be used to detect anomolies and reveal information about the underlying traffic [41].

4. Protecting stored data

4.1. Stored e-mail

Some secure mail systems store original message ciphertext and decrypt only for display. While this protects messages on disk, it means that keys must be stored until all messages they protect are deleted. We must assume that an attacker has copies of message ciphertext sent over an insecure network such as the Internet. These messages remain vulnerable until the corresponding private key is deleted.

Messages may be stored temporarily encrypted with a short-lifetime key, but only for the key's lifetime; they are unreadable once it has been deleted. They should be stored encrypted under a long-term storage key. Mail clients may implement their own secure storage facilities, or use those provided by other software.

4.2. Steganographic File Systems

Encrypted filesystems protect stored data against access without the appropriate key. Good system design should protect that key against access even by powerful adversaries. But it cannot stop its compelled disclosure by a user. Coercion has long been recognised as a threat to security systems, whether by banks protecting vaults against kidnapped managers or intelligence agencies providing plausible deniability for their agents.

Steganographic filesystems can protect against such compulsion. They allow each file to be protected by a password. Only someone who knows a file's name and password can even tell that the file exists [1]. While the fact that a steganographic filesystem is present can be detected, a user need simply reveal the password of a small number of rela-

tively innocuous files. The existance of any other protected information simply cannot be shown.

Anderson's two original designs used either linear algebraic operations or ciphers to disguise files. The former creates a set of intially random cover files that are modified so that stored data can be retrieved by XOR'ing some subset described by the filename and password. The number of files must be large enough to prevent an attacker trying all possible subsets. This system provides provable security without requiring the use of a strong cipher, but its read and write operations are slow, and known plaintext causes vulnerability. This is a weakness with mass-market software that uses many standard file headers and formats. The second system hides encrypted blocks in a large amount of random data at locations derived from the filename and password. As collisions will occur, particularly once more than $\sqrt{n}$ of the n blocks of the filesystem are used (due to the birthday paradox), redundant copies are also created.

An implementation for Linux has provided the facility for storing a number of files at one of up to 15 security levels, each protected by a different passphrase. Only the files stored at a security level equal to or less than that currently opened by a user are visible [28]. This is implemented using a block allocation table where each file's blocks are listed, encrypted using the relevant security level's key. The table itself is steganographically protected. Data can be hidden in the unused blocks of another filesystem, removing Anderson's requirement for a separate data partition.

4.3. Secure remote storage

Some of the threats we have described are location-dependent. Laptops and their users, for example, are at higher risk in Bogotá than London.

With global Internet connectivity, there is little reason why data should be put at risk by physical transport through these areas. At the very least, users should move sensitive information from laptop disks to Internet-accessible secure storage before they travel overseas with their machines. Ephemeral data such as the content of Web caches or operating system swapfiles should be securely wiped. Data can then be accessed at the user's destination over a secure Internet connection. Users may simply re-synchronise their files at that point, or mount a remote filesystem.

Users should be particularly careful not to expose cryptographic keys or passwords to disclosure. These may be cached by applications or 'single sign-on' systems, or in swapfiles.

4.4. Storing keys in tamperproof hardware

Cryptographic keys are are the most valuable of information assets. Their disclosure can lead to the compromise of system, data and communications security until the keys are revoked or expire.

Tamper-evident smartcards are therefore excellent key storage devices. They will resist most attacks that attempt to gain access to key material, as keys never legally leave the card but are used by its processor to directly cryptographically process data. A successful attack should require physical access to the card, and leave obvious evidence, so that keys can be revoked when their owners realise a card has been stolen or damaged. Measures such as randomised clock signals, structural destruction of test circuitry and top-layer sensor meshes have been suggested as ways to provide this smartcard protection. Unfortunately, few cards yet provide adequate security [25].

5. Procedural safeguards

5.1. Secret splitting

Multinational businesses can minimise their exposure to particular jurisdictions by splitting information cryptographically between them. A rarely used, high value key – such as a company's root certification key – could require the approval of n of m corporate officers in different countries for access. Each should be required to verify that the others are not operating under compulsion before giving assent.

Company security functions should also be judiciously sited. Any sensitive key material should be kept out of the reach of key warrants in a country like Ireland that has explicitly rejected them, and in areas with effective law-enforcement protection against criminals.

5.2. Designated revokers

All well-designed public key infrastructures allow keys to be revoked. Revocation lists should always be checked before keys are used. Once a key has been disclosed, it is vital to revoke it. However, this may not be practically possible, because at the time of seizure the key holder may be incapacitated. It is therefore important that keyholders should designate third parties to revoke their keys when required to do so. Procedures should be in place to notify these third parties if the key holder is out of contact for a specified time period, or some other event occurs that may signify their potential compromise.

Designated revokers normally use one of two methods: they either have a pre-prepared revocation, signed by the key that will be revoked, or the PKI has the capability of designating revocation keys other than the key to be revoked.

Ideally, software should only allow the disclosure of revoked keys. This will prevent any further traffic being vulnerable to decryption using that key. The revocation should

be available to anyone who may attempt to use that key in future.

5.3. Backups

It is vital that short lifetime private keys *not* be backed up. This can be surprisingly difficult to achieve, particularly in corporate environments, and needs to be regularly checked. It is likely to require special exceptions to standard backup procedures, which may be accidentally lost when equipment and software is upgraded. If keys are lost, they will quickly expire, but can anyway be revoked and replaced.

Frequent audits of key storage and its relationship to backup procedures are recommended to avoid inadvertant long-term storage of private keys.

Cryptographic techniques may be used to minimise the data excluded from backup. For example, a set of short-lifetime private keys $k_{1\ldots n}$ can be stored with k_i encrypted symetrically using $hash(k_{i-1})$. Only the current key needs protection from backup. Just before wiping, it should be used to decrypt the next key in the series. This technique can be particularly useful with limited storage environments such as smartcards.

5.4. Document destruction

Like keys, many company documents should have strictly limited lifetimes.

Regulators require certain classes of information to be kept for varying periods of time. The US Securities and Exchange Commission, for example, has ruled that brokers must keep transaction and participant information for automated trading systems for three years [37]. Lawyers suggest that otherwise documents should generally be kept one year beyond the local relevant statute of limitations [8].

There is often little reason why documents should be kept for one day longer. Companies should have procedures in place to allow the effective destruction of all copies of a document, including backups. These policies should be followed as strictly as possible, as selective document destruction may lead investigators to assume specific incriminating information has been purged [40].

E-mail destruction policies are especially important. Because users tend to regard messages as closer to telephone conversations than letters, they metaphorically commit to paper information that should never have been given permanence. Copies of indiscrete messages on disks, backup tapes, or mail servers can threaten the very existence of a company, as the Microsoft case has shown. An e-mail discovery request can also prove extremely expensive for companies with large collections of old messages. Computer Forensics, Inc. estimate that "the process of reviewing and collecting e-mail from disparate tapes and databases usually costs at least $25,000" [19].

Companies may wish to arrange for e-mail to be transferred as directly from sender to recipient, over forward secure links, as possible. This will minimise the number of intermediate points at which messages are stored.

If litigation is commenced against a company, it must immediately take steps to prevent the destruction of relevant information. However, limiting the existence of copies of a document will allow it to be promptly disposed of once the matter has been resolved.

6. Conclusion

We are rushing headlong into a global information economy. Information will continue to increase in value as the most precious of corporate resources. Companies must protect these assets as zealously as that valuation demands, against even extreme threats such as organised crime or governments performing economic espionage.

We have described relatively simple enhancements to security services that are a first step towards providing this protection. All communications should use forward secret links that minimise traffic data. Stored data should be kept safe on steganographic filesystems, and have a strictly limited lifetime. Information should be split between and sited in relatively trusted jurisdictions, with damage limitation procedures in place in case of compromise.

These measures will protect company information *and* personnel. A user who has access only to the minimum set of information required is a less valuable target for criminals, decreasing the likelihood of blackmail or physical attack upon them.

None of this is to suggest that companies should break the law. When choosing to do business in a country, organisations must respect the obligations that come with that choice. But these do not include making more information than is required available to third parties. The measures suggested in this paper are legal in the vast majority of the world's countries.

In the long term, we must hope governments will also respect the spirit if not the letter of international law, and realise that it is in all of our shared interest that they cease undertaking economic espionage. Brian Gladman has compared the current situation to the encouragement of piracy by the maritime nations in the fifteenth and sixteenth centuries. "Eventually, however, it was recognised that this was undermining the development of world trade and was not in the real interests of any of the nations involved in it. In consequence state sponsorship ceased and the rule of law was established on the high seas. At the moment we have some of the 'large' nations of the world acting as pirates on the information infrastructure and we are trying at the same time

to develop this infrastructure as a basis for electronic commerce and the information society. As happened a few centuries ago, nations will have to stop sponsoring 'information piracy' if we are ever to move forward to the global information society that we all hear so much about. This change is inevitable but it is not clear how long it will take." [20]

Until we reach this point, the companies whose information assets are at risk will continue to need the kind of extreme defensive mechanisms we have outlined.

7. Acknowledgements

Thanks to Ross Anderson, Adam Back, Nicholas Bohm, Richard Clayton, Brian Gladman, Peter Gutmann, William Murray and the anonymous referees, whose useful comments have have been incorporated into this paper.

References

[1] Ross Anderson, Roger Needham and Adi Shamir. The Steganographic File System. In Proc. Second International Information Hiding Workshop, Portland, Oregon, USA, April 1998. http://www.cl.cam.ac.uk/ftp/users/rja14/sfs3.ps.gz

[2] R. Atkinson. Security Architecture for the Internet Protocol. RFC 1825, August 1995.

[3] James Bamford. The Puzzle Palace. New York: Penguin Books, 1983, pp. 126–7.

[4] David Banisar and Wayne Madsen. Cryptography and Liberty 2000: An International Survey of Encryption Policy. Washington, DC: EPIC, March 2000. http://www2.epic.org/reports/crypto2000/overview.html#Heading10

[5] Hale v. Henkel, 201 U.S. 43, 74–75, March 1906.

[6] Duncan Campbell. Dispatches: The Hill. Channel 4 Television (UK), 6 October 1993.

[7] Duncan Campbell. Development of surveillance technology and risk of abuse of economic information. Report of the European Parliament Scientific and Technical Options Assessment Panel, April 1999. http://www.gn.apc.org/duncan/stoa.htm

[8] Ken Carroll. Document Retention/Destruction Policies. Miller & Martin Advisor, September 1998. http://www.millermartin.com/pubs/adv12_98_docret.htm

[9] CCITT. Recommendation X.509: The Directory — Authentication Framework, 1988.

[10] D. Chaum. Untraceable Electronic Mail, Return Addresses, and Digital Pseudonyms. Communications of the ACM, 24(2) 84–88, February 1981.

[11] Council of Europe. Draft Convention on Cyber-crime, April 2000. http://conventions.coe.int/treaty/en/projets/cybercrime.htm

[12] Simon Davies. Travelers face jail if they refuse to let Customs scan their laptops. The Daily Telegraph, 16 September 1998. http://www.sightings.com/political/laptops.htm

[13] Whitfield Diffie and Martin Hellman. New directions in cryptography. IEEE Transactions on Information Theory 22(6), November 1976, 644-654.

[14] US Department of Defense. Department of Defense Trusted Computer System Evaluation Criteria. DoD 5200.28-STD, December 1985.

[15] Ireland Electronic Commerce Act 2000, 6 April 2000. http://www.entemp.ie/ecd/ecb12000.pdf

[16] T. Elgamal. A Public Key Cryptosystem and a Signature Scheme Based on Discrete Logarithms. IEEE Transactions on Information Theory 31(4), July 1985, 469–472.

[17] Executive Order 12333 — United States intelligence activities. 46 FR 59941, 3 CFR, December 1981. http://www.nara.gov/fedreg/eos/e12333.html

[18] Ned Freed. Encrypting RFC822 headers in S/MIME or PGP/MIME messages. Post to ietf-open-pgp@imc.org, September 1998. http://www.imc.org/ietf-openpgp/mail-archive/msg01941.html

[19] Roberta Fusaro. Cases highlight need for e-mail policies. ComputerWorld, 10 May 1998. http://www.computerworld.com/home/print.nsf/(frames)/9810056D2A?OpenDocument&~f

[20] Brian Gladman. DTI to ban electronic export of crypto from the UK! Post to ukcrypto@maillist.ox.ac.uk, July 1998. http://www.cs.ucl.ac.uk/staff/I.Brown/archives/ukcrypto/0898-1198/msg00060.html

[21] Brian Gladman. The Regulation of Investigatory Powers Bill — The Provisions for Government Access to Keys. Foundation for Information Policy Research report, February 2000. http://www.fipr.org/rip/RIPGAKBG.pdf

[22] David M. Goldschlag, Michael G. Reed, and Paul F. Syverson. Onion Routing for Anonymous and Private Internet Connections. Communications of the ACM 42(2), February 1999. http://www.onion-router.net/Publications/CACM-1999.ps

[23] Dan Goodin and Jeff Pelline. Smoking gun in Microsoft memos? CNET News.com, May 1998. http://news.cnet.com/news/0-1003-200-329400.html

[24] UK Intelligence Services Act, May 1994.

[25] O. Kömmerling and M. G. Kuhn. Design Principles for Tamper-Resistant Smartcard Processors. Proc. USENIX Workshop on Smartcard Technology, May 1999.

[26] Courtney Macavinta. Yahoo message board suit continues. CNET News.com, March 1, 1999. http://news.cnet.com/news/0-1005-200-339276.html

[27] Declan McCullagh. Helsingius shuts down anon.penet.fi server in Finland. Politech, August 1996. http://www.itu.reading.ac.uk/misc/mailing_lists/rre/00000167.htm

[28] Andrew D. McDonald and Markus G. Kuhn. StegFS: A Steganographic File System for Linux. Proc. Third International Workshop on Information Hiding, Dresden, Germany, October 1999. http://www.cl.cam.ac.uk/~mgk25/ih99-stegfs.pdf

[29] W. H. Murray. UK Customs Check for laptop porn. Posts to talk.politics.crypto, August 1998. http://privacy.nb.ca/cryptography/archives/cryptography/html/1998-08/0111.html

[30] Chris Nuttall. UK Customs check for laptop porn. BBC News Online, 13 August 1998. http://news.bbc.co.uk/hi/english/sci/tech/newsid_150000/150465.stm

[31] OECD Guidelines for Cryptography Policy, December 1997. http://www.oecd.org/dsti/sti/it/secur/prod/crypto2.htm

[32] John Pike. Intelligence Agency Budget and Personnel. Federation of American Scientists, July 1998. http://www.fas.org/irp/agency/budget1.htm

[33] UK Regulation of Investigatory Powers Act, July 2000. http://www.legislation.hmso.gov.uk/acts/acts2000/10000023.htm

[34] R. Rivest, A. Shamir and L. Adleman. A Method for Obtaining Digital Signatures and Public-Key Cryptosystems. Communications of the ACM, 21(2) 120-126, February 1978.

[35] Nicholas Rufford. MI5 builds new centre to read e-mails on the net. The Sunday Times, 30 April 2000. http://www.the-times.co.uk/news/pages/sti/2000/04/30/stinwenws01034.html

[36] B. Schneier and C. Hall. An Improved E-mail Security Protocol. 13th Annual Computer Security Applications Conference. New York: ACM Press, 1997, pp. 232–238.

[37] Securities and Exchange Commission. 17 CFR 240, 12 June 1995. http://www.sec.gov/rules/final/17a23.txt

[38] Greg S. Sergienko. Self Incrimination and Cryptographic Keys. The Richmond Journal of Law and Technology 1, February 1996. http://www.richmond.edu/jolt/v2i1/sergienko.html

[39] Scott Shane. Mixing business with spying; secret information is passed routinely to US. Baltimore Sun, 1 November 1996.

[40] Peter L. Simmons. Records Retention Policies: When Is It Safe To Destroy Documents. Fried, Frank, Harris, Shriver & Jacobson client memos, January 1997. http://www.ffhsj.com/firmpage/cmemos/131280.htm

[41] Brenda Timmerman. Secure dynamic adaptive traffic masking. Proc. New Security Paradigms Workshop, Caledon Hills Canada, September 1999.

[42] Mark Urban. UK Eyes Alpha. London: Faber and Faber, 1996, p.235.

[43] Robert Windrem. U.S. steps up commercial spying. MSNBC News, 7 May 2000. http://msnbc.com/news/403435.asp?cp1=1

[44] John Woolsey. Foreign Press Center, Washington, DC briefing, 7 March 2000. http://cryptome.org/echelon-cia.htm

[45] P. R. Zimmerman. The Official PGP Users Guide. Boston: MIT Press, June 1995.